arXiv:1604.08189v1 [math.OC] 27 Apr 2016

Integrating Storage to Power System Management

Luckny Zéphyr, C. Lindsay Anderson
Biological and Environmental Engineering Department, Cornell University

April 28, 2016

Abstract

Wind integration in power grids is very difficult, essentially because of the uncertain nature of wind speed. Forecasting errors on output from wind turbines may have costly consequences. For instance, power might be bought at highest price to meet the load. On the other hand, in case of surplus, power may be wasted. Energy storage facility may provide some recourse against the uncertainty on wind generation. Because of the sequential nature of power scheduling problems, stochastic dynamic programming is often used as solution method. However, this scheme is limited to very small networks by the so-called curse of dimensionality. To face such limitations, several approximate approaches have been proposed. We analyze the management of a network composed of conventional power units as well as wind turbines through approximate dynamic programming. We consider a general power network model with ramping constraints on the conventional generators. We use generalized linear programming techniques to linearize the problems. We test the algorithm on several networks of different sizes and report results about the computation time. We also carry out comparisons with classical dynamic programming on a small network. The results show the algorithm seems to offer a fair trade-off between solution time and accuracy.

Key words: Power grid management; Energy storage; Stochastic Dynamic Programming; Stochastic Dual Dynamic Programming; Approximate dynamic Programming; Generalized linear programming.

1. Introduction

Cost-effective management of power units is a very challenging task. Generating units have to be committed such that demand for electricity is met. This is more difficult to carry out because of significant load variability, among others.

Commitment of traditional units such as nuclear and coal-fire involve cost, mostly because of fuel consumption, and is source of environment concerns. Driven by both increasing environmental awareness and technological advances, since the last decade, wind-based electricity generation has been widely promoted [63]. For instance, the 2001/77/EC European Commission Directive had set at 22% the renewable integration target for Europe by 2010

[35]. In the United States, in 2006, the total wind installed capacity grew approximately from 9 000 MW to 11 600 MW [52]. As a result of political emphasis, in 2006, on the need to increase the United States energy efficiency and to diversify the energy portfolio, a collaborative initiative was created to explore the requirements for a 20% wind integration scenario by 2030 [30].

However, unlike conventional power source, output from wind turbines is uncontrollable. Due to the random nature of such power, these latter cannot serve both supply and reliability purposes (in case of outage in the network). As a consequence, wind generators cannot replace conventional ones to meet peak load [35]. The intermittency of wind generation may create an imbalance between the supply and demand for power. In case of excess generation, power may be curtailed due to transmission congestion, and in case of peak load, to meet the demand, power would probably be bough at expensive price because of inadequate anticipation (forecasting error).

Energy storage devices may serve as recourse to circumvent the uncertainty of wind power. Such devices may be used to store excess generation, or for arbitrage profits. Indeed power might be bough at a lower price during off-peak hours to be stored and sold at a higher price during peak load hours [33]. For an in depth description of existing and in development storage technology, see [21, 19]. Benefits as well as market-related questions are discussed in [15].

A significant stream of research has been focusing on harvesting wind power in the presence of storage. The benefits of coupling intermittent renewable energy with storage are discussed in [5, 31, 6, 13]. Prior studies have also analyzed the coupling of wind generation with storage via stochastic programming (e.g. [17], [1], [32]). Unlike other research where a fixed wind generation curve is used, [54] analyze the optimal wind turbine rating and the storage configuration of a wind farm coupled to compressed air energy storage. [48] compare approximate dynamic programming schemes for energy storage management based on instrumental Variables and projected Bellman errors for a model ignoring transmission lines in the presence of a single storage device. In [63], the valuation of a single storage unit provided by a stochastic approach is compared with several valuation heuristics for a model with limited transmission capacity. Storage of energy for arbitrage opportunities is studied by [33] via stochastic programming and dynamic programming.

We examine the management of a power network composed of conventional units, wind turbines, as well as storage devices. In particular, we discuss an approximate stochastic method based on stochastic dual dynamic programming for the management of the storage units. Section 2 deals with the operation of a power network equipped with energy storage facilities. Section 3 provides a formulation of the problem under the framework of stochastic dynamic programming and discusses the associated sources of complexity. Section 4 discusses a linear approximation of the problem. In Section 5, we review some general approximate dynamic programming techniques widely used in power system management, while Section 6 presents an approximate formulation for the problem. We illustrates the algorithm's performance over a small network and carry out comparisons with dynamic programming in Section 7. Results of additional experiments on larger networks are reported in Section 8.

2. Operation and control of a power system with storage

Optimization models are often used to determine optimal *generation schedule* for the conventional generators over a finite planning horizon, T, in order to meet the demand in each time step t at minimal cost. This is usually achieved by formulating the problem as a two-stage optimization problem, where in the first stage decisions are made about on/off status of the conventional generators based on anticipation (or forecasting) on demand and wind generator output. The second stage decisions are usually concerned with generation (*economic dispatch* of the committed generators), as well as recourse actions (reserve requirements) to offset forecasting errors and for reliability purposes in case of unforeseen events such as generator failure, transmission line disruption. The time step is usually an hour or a fifteen-minute interval, and the planning span is usually 24, 48 or 168 hours (a week).

Based on load and wind generation forecasts, the operator makes decision about the commitment of the conventional generators in each time period of the horizon. We assume that generation decisions (for conventional generators) are made after observations of the random variables (wind generator output).

Here, we are dealing with the economic dispatch problem alone, assuming that the conventional generators are all committed. For instance, [4] first use a simulator to determine the on/off status of the traditional generators in each point of time, then solve the dispatch problems.

A power system may be represented by a directed graph $\mathcal{G} = (N, L)$, where each node $n \in N$ represents a bus, and each link a transmission line. In each node may be located a set of conventional generators G_n, a wind farm, and a storage device. We distinguish the different components of the graph via the following set of indices: (i) G is the set of conventional generators, (ii) $M \subseteq N$ is the set of wind farms, and (iii) S denotes the set of storage devices.

The generators are mechanical devices, they have to operate within finite generating limits. While in motion, a minimal output may be required for a generator to be in steady state. Similarly, a threshold is imposed on the maximal output to avoid mechanical damages. For any generator $g \in G, \underline{p}_g$ and $\bar{p}_g$ (in MW) denote these bounds, respectively. In addition, there usually are minimum and maximum allowable changes in the generations between two consecutive periods. For any generator $g \in G$, the following defines ramp down and ramp up limits:

$$-\underline{\lambda}_g \leq p_{gt} - p_{g,t-1} \leq \bar{\lambda}_g.$$

In addition to power generated in a node, power can flow between two nodes trough transmission lines defined as the set of undirected pairs $L \subseteq N \times N$ (assuming power can flow in either direction). For any node $n \in N$, define $O_n = \{(n, j) \in L\}$ to be the set of transmission lines that leave node n. Similarly, let $I_n = \{(j, n) \in L\}$ be the set of transmission lines that enter node n.

For any node n, $\hat{D}_{nt}$ denotes a "reliable" forecast (in MW) of the load (demand), in period t. This demand is supplied by the generation in the node, by power transported to the node through the transmission lines, or by discharging the storage, if any, or by some combination of power from the three. Thus, p_{gt} denotes power generation, in MW, from generator g in period t at the cost $CP_{gt}(p_{gt})$.

The output from the wind turbines is random. Indeed, generation from such units is conditioned on wind speed, which depends on uncontrollable meteorological conditions. Therefore, $\tilde{W}_t = (\tilde{W}_{1t}, \ldots, \tilde{W}_{|M|t})$ denotes the random vector of outputs from the wind turbines in period t; w_{mt}, in MW, is a particular realization of the stochastic process $\{\tilde{W}_{mt}\}$.

We note by e_{lt} the power, in MW, flowing through transmission line l in period t. The transmission lines have limited capacity, thus upper bounds are imposed on the power flowing through to prevent disruption. $\bar{e}_l, l \in L$, denote such bounds. We assume bidirectional lines, i.e., power can flow in either direction of the line. Therefore, for any transmission line $l = (n, n') \in L$, in period t, if e_{lt} is positive, power is transmitted from node n to node n'. A negative value indicates the opposite. The power flowing through a transmission line is proportional to the difference between the phase angles, in radians, of the two end buses, i.e., $e_{lt} = B_l(\theta_{n't} - \theta_{nt})$, where B_l is the susceptance of line l [38]. To avoid that the system is over-determined, i.e. with more equations than unknown, the voltage angle at the reference (slack) bus is usually set to zero.

In each period t, in addition to generation decisions, the operator of the network also makes decisions on the use of the storage. In each node where a storage device is located, in each period, the device may be charged or discharged to compensate for imbalance between the supply and demand. Such imbalance is more likely to occur in peak demand hours, i.e., periods where electricity consumption is highest. For instance, demand for power tends to be lower at night when people are asleep. During the summer, peak demand tends to occur in the afternoon because of higher activity rate and use of air conditioning by buildings. In the winter, demand is often higher in the morning and in the evening due to greater need for heating. Operators usually prepare for peak demand by committing extra power plants that may be called upon quickly in periods of higher demand or failures (reserve).

For any storage device $j \in S$, define s_{jt} to be its level of energy (MWh) in the beginning of period t (or the end of period $t-1$). This energy is converted into power (MW) via the simple equation Energy = Power $\times$ Time. We note by Δ_{jt} the variation in the level of charge from the beginning through the end of period t. A negative value of Δ_{jt} indicates discharging, and a positive value the opposite. There may be a cost associated with varying stored energy level. $CS_{jt}(\Delta_{jt})$ is the cost associated with such variation, if any. Such cost is proportional to the amount of energy stored or discharged [47].

The storage units have limited capacity, and in order to last, they cannot be completely depleted. Consequently, the level of charge can only varied progressively over time. We therefore note by $\underline{s}_j$ and $\bar{s}_j$ lower and upper bounds on the level of storage of device j. Assume that the discharge and charge maximum capacity is the same, in period t, the variation in the level of device j then obeys

$$-\underline{\Delta}_j \leq \Delta_{jt} \leq \bar{\Delta}_j.$$

Lastly, each storage facility has input and output efficiency; energy is lost both at charging and discharging. We then define c_j (resp. d_j) to be the efficiency coefficient of charging (resp. discharging) of storage device $j \in S$, where $0 < c_j \leq 1$, and $0 < d_j \leq 1$.

Imbalance between the load and supply may result because of wind power forecasting errors. Since the storage units have limited capacity, they may not be able to absorb the total production surplus (charging) or to deliver the difference (discharging) in case of power

shortage. For any node $n \in N$, $\gamma_{nt} \in \mathbf{R}$ denotes the power absorbed or delivered by the storage unit. In case of imbalance, we assume that power excess (shortage) is absorbed (delivered) at a high rate. For each node $n \in N, \kappa^+_{nt} \in \mathbf{R}^+$ ($\kappa^-_{nt} \in \mathbf{R}^+$) denote such excess (shortage); $\mathbf{R}$ and $\mathbf{R}^+$ are the set of real and non-negative real numbers, respectively.

Let us also explicitly distinguish charging and discharging via the variables $\Delta^+_{jt} \in \mathbf{R}^+$, and $\Delta^-_{jt} \in \mathbf{R}^+, j \in S$, respectively. Thus, in each period t, we have the power balance equations

$$\sum_{k=1}^{N_n} p_{kt} + \sum_{l \in I_n} e_{lt} + w_{nt} - \delta_n \gamma_{nt} - \kappa^+_{nt} + \kappa^-_{nt} = \sum_{l \in O_n} e_{lt} + \hat{D}_{nt}, \forall n \in N,$$

with

$$\delta_n = \begin{cases} 1 & \text{if there is a storage facility at node n,} \\ 0 & \text{otherwise.} \end{cases}$$

By convention, $w_{nt} = 0, n \in N$, if there is no wind farm at node n. We also have

$$\Delta^+_{jt} = \max\{0, c_j \gamma_{jt}\}, j \in S,$$

and

$$\Delta^-_{jt} = \max\left\{0, -\frac{\gamma_{jt}}{d_j}\right\}, j \in S.$$

The storage level then evolves according to

$$s_{j,t+1} = \alpha_j s_{jt} + \Delta^+_{jt} - \Delta^-_{jt}, \forall j \in S,$$

where $\alpha_j \in (0, 1]$ is the storage efficiency of device j. As defined, it is obvious that both charging and discharging will not occur at the same time.

We aim to find the *policy* $(P^*, \Delta^*) = [(P^*_1, \Delta^*_1), \cdots, (P^*_T, \Delta^*_T)]$ that minimizes the expected operating cost over the entire planning horizon, where $P^*_t = (p^*_{1t}, \cdots, p^*_{|G|t})$, and $\Delta^*_t = (\Delta^*_{1t}, \cdots, \Delta^*_{|S|t}), 1 \le t \le T$; Δ_t refers to either charging or discharging. (P^*, Δ^*) then solves

$$\min_{P_t,\Delta_t,\kappa_t,s_{t+1}} \left\{ \mathbb{E}\left[\sum_{t=1}^{T} \left(\sum_{g=1}^{|G|} CP_{gt}(p_{gt}) + \sum_{j=1}^{|S|} CS_{jt}(\Delta_{jt}) + M \sum_{n=1}^{N} (\kappa_{nt}^{+} + \kappa_{nt}^{-}) \right) \right] \right\} \tag{1}$$

Subject to, for $1 \leq t \leq T$:

$$\sum_{k=1}^{N_n} p_{kt} + \sum_{l \in I_n} e_{lt} + w_{nt} - \delta_n \gamma_{nt} - \kappa_{nt}^{+} + \kappa_{nt}^{-} = \sum_{l \in O_n} e_{lt} + \hat{D}_{nt}, \qquad n \in N \tag{2}$$

$$p_{gt} - p_{g,t-1} \leq \bar{\lambda}_g, \qquad g \in G \tag{3}$$

$$p_{g,t-1} - p_{gt} \leq \underline{\lambda}_g, \qquad g \in G \tag{4}$$

$$\underline{p}_g \leq p_{gt} \leq \bar{p}_g, \qquad g \in G \tag{5}$$

$$e_{lt} = B_l(\theta_{n't} - \theta_{nt}), \qquad l = (n, n') \in L \tag{6}$$

$$-\bar{e}_l \leq e_{lt} \leq \bar{e}_l, \qquad l \in L \tag{7}$$

$$s_{j,t+1} = \alpha_j s_{jt} + \Delta_{jt}^{+} - \Delta_{jt}^{-}, \qquad j \in S \tag{8}$$

$$\underline{s}_j \leq s_{j,t+1} \leq \bar{s}_j, \qquad j \in S \tag{9}$$

$$0 \leq \Delta_{jt}^{+} \leq \bar{\Delta}_j, \qquad j \in S \tag{10}$$

$$0 \leq \Delta_{jt}^{-} \leq \bar{\Delta}_j, \qquad j \in S \tag{11}$$

$$\gamma_{jt} = \frac{\Delta_{jt}^{+}}{c_j} - d_j \Delta_{jt}^{-}, \qquad j \in S \tag{12}$$

$$\kappa_{jt}^{+} \geq 0, \qquad j \in S \tag{13}$$

$$\kappa_{jt}^{-} \geq 0, \qquad j \in S \tag{14}$$

$$\Delta_{jt}^{+} \geq 0, \qquad j \in S \tag{15}$$

$$\Delta_{jt}^{-} \geq 0, \qquad j \in S \tag{16}$$

$\mathbb{E}$ is the expectation operator, which is taken over $\tilde{W}_t$. For any set X, |X| denotes its cardinality. M is a big number.

The multi-period problem (1–16) can theoretically be solved to optimality if functions CP_{gt} is convex in p_{gt}, CS_{jt} being proportional to Δ_{jt}. Indeed, the set $\Psi_t = \left\{ (P_t, e_t, \tilde{W}_t, \theta_t, s_t, \Delta_t^{+}, \Delta_t^{-}, \gamma_t, \kappa_t^{+}, \kappa_t^{-}) | (2-16) \right\}$ is a polyhedron. Therefore, by convexity of the cost functions, (1–16) is a convex problem. However, the problem may not be tractable numerically if we want a detailed representation for the underlying process of $\tilde{W}_t$.

3. Representation under the framework of stochastic dynamic programming

Problem (1–16) is a multi-period stochastic program. It also is a sequential decision problem. In fact, in each period t, the operator of the system observes the level of the stored energy as well as the generations in the previous period, and based on updated forecast for the wind turbines outputs (or wind speed) and the observation of the previous outputs, he/she

decides about the generation of the conventional generators as well as the variation in the storage (charging or discharging) in order to meet the demand in each node. Thus, the tuple (s_t, p_{t-1}, w_{t-1}) will be called the *state of the system*, or state for short.

Let us observe that whereas p_{t-1} is known, we see from constraints (3-5) that $p_{g,t-1} - \underline{\lambda}_g$ and $\underline{p}_g$ are two minorants for $p_{gt}, g \in G$. Similarly, $p_{g,t-1} + \bar{\lambda}_g$ and $\bar{p}_g$ are two majorants for p_{gt}. Consequently, we have

$$\max\{p_{g,t-1} - \underline{\lambda}_g, \underline{p}_g\} \leq p_{gt} \leq \min\{p_{g,t-1} + \bar{\lambda}_g, \bar{p}_g\}.$$

In principle, stochastic dynamic programming (SDP) is suited for problem (1–16). SDP sequentially decomposes (by period) the overall problem into smaller subproblems in a coordinated way, by seeking the best trade-off between the immediate and future use of the storage. Define $F_t(s_t, p_{t-1}, w_{t-1})$ to be the *cost-to-go function* from the beginning of period t to the end of the horizon. In addition, assume that the process $\{\tilde{W}_t\}$ is Markovian, i.e., $\mathbb{P}(\tilde{W}_t = w_t | \tilde{W}_{t-1} = w_{t-1}, \ldots, \tilde{W}_0 = w_0) = \mathbb{P}(\tilde{W}_t = w_t | \tilde{W}_{t-1} = w_{t-1})$. Therefore, for $t = T, T-1, \cdots, 1$, a SDP recursion associated with problem (1–16) is given by

$$F_t(s_t, p_{t-1}, w_{t-1}) = \min_{P_t, \Delta_t, \kappa_t, s_{t+1}} \left\{ \sum_{g=1}^{|G|} CP_{gt}(p_{gt}) + \sum_{j=1}^{|S|} CS_{jt}(\Delta_{jt}) + M \sum_{n=1}^{N} (\kappa_{nt}^+ + \kappa_{nt}^-) + \right.$$
$$\left. \mathbb{E}_{\tilde{W}_t | w_{t-1}} \left[F_{t+1}(s_{t+1}, p_t, \tilde{W}_t) \right] \right\} \tag{17}$$

$$\text{S.t.} \sum_{k=1}^{N_n} p_{kt} + \sum_{l \in I_n} e_{lt} + w_{nt} - \delta_n \gamma_{nt} - \kappa_{nt}^+ + \kappa_{nt}^- = \sum_{l \in O_n} e_{lt} + \hat{D}_{nt}, \ n \in N \tag{18}$$

$$p_{gt} \geq \max\{p_{g,t-1} - \underline{\lambda}_g, \underline{p}_g\}, \qquad g \in G \tag{19}$$

$$p_{gt} \leq \min\{p_{g,t-1} + \bar{\lambda}_g, \bar{p}_g\}, \qquad g \in G \tag{20}$$

$$e_{lt} = B_l(\theta_{n't} - \theta_{nt}), \qquad l = (n, n') \in L \tag{21}$$

$$-\bar{e}_l \leq e_{lt} \leq \bar{e}_l, \qquad l \in L \tag{22}$$

$$s_{j,t+1} = \alpha_j s_{jt} + \Delta_{jt}^+ - \Delta_{jt}^-, \qquad j \in S \tag{23}$$

$$\underline{s}_j \leq s_{j,t+1} \leq \bar{s}_j, \qquad j \in S \tag{24}$$

$$0 \leq \Delta_{jt}^+ \leq \bar{\Delta}_j, \qquad j \in S \tag{25}$$

$$0 \leq \Delta_{jt}^- \leq \bar{\Delta}_j, \qquad j \in S \tag{26}$$

$$\gamma_{jt} = \frac{\Delta_{jt}^+}{c_j} - d_j \Delta_{jt}^-, \qquad j \in S \tag{27}$$

$$\kappa_{jt}^+ \geq 0, \qquad j \in S \tag{28}$$

$$\kappa_{jt}^- \geq 0, \qquad j \in S \tag{29}$$

$$\Delta_{jt}^+ \geq 0, \qquad j \in S \tag{30}$$

$$\Delta_{jt}^- \geq 0, \qquad j \in S \tag{31}$$

The complexity of problem (17–31) essentially stems from two fronts, namely (i) computing an expectation, and (ii) finding an optimal policy $\Pi^*(s_t, p_{t-1}, w_{t-1})$. The complexity

of (i) is related to the dimension of the random vector $\tilde{W}_{t-1}$, and step (ii) may be prohibitive because of the dimension of the joint state space $(s_t, p_{t-1}, \tilde{W}_{t-1})$, except for very rare cases, such as *linear systems with quadratic costs* where the optimal policy is a linear function of the state variables (see [11, pp. 148-149]).

4. Approximation through linear programming

We will later implement a stochastic dynamic programming scheme to approximate the cost-to-go function. This method assumes linear cost functions, which, however, are usually non-linear. Non-linearity issues may be circumvented via inner *generalized linear programming* (GLP) (see [51]). This is a technique to perform interpolations over a sample of points. In the case of convexity, GLP converges to the original problem for well conceived sample of points. In case of non-convexity, a convex approximation is performed for the original problem. For each generator $g \in G$, suppose we have exact evaluations of the cost function over the sample of points $\{\hat{p}_{gj} | j \in \Lambda_g\}$. Assume that $\mathbb{E}_{\tilde{W}_t|w_{t-1}}\left[\tilde{F}_{t+1}(s_{t+1}, p_t, \tilde{W}_t)\right]$ is a linear function (we will later substitute the expectation with a single decision variable), the following is a linear approximation to the original problem (17–31):

$$\tilde{F}_t(s_t, p_{t-1}, w_{t-1}) = \min_{P_t, \Delta_t, \kappa_t, s_{t+1}, \beta_t} \left\{ \sum_{g=1}^{|G|} \sum_{j \in \Lambda_g} \beta_{gj,t} CP_g(\hat{p}_{gj}) + \sum_{j=1}^{|S|} CS_{jt}(\Delta_{jt}) + M \sum_{n=1}^{N} (\kappa_{nt}^{+} + \kappa_{nt}^{-}) + \right.$$
$$\left. \mathbb{E}_{\tilde{W}_t|w_{t-1}}\left[\tilde{F}_{t+1}(s_{t+1}, p_t, \tilde{W}_t)\right] \right\} \tag{32}$$

$$\text{S.t.} (17) - (31) \tag{33}$$

$$p_{gt} = \sum_{j \in \Lambda_g} \beta_{gj,t} \hat{p}_{gj}, \qquad g \in G \tag{34}$$

$$\sum_{j \in \Lambda_g} \beta_{gj,t} = 1, \qquad g \in G \tag{35}$$

$$\beta_{gj,t} \geq 0, \qquad g \in G, j \in \Lambda_g \tag{36}$$

In each period t, for each generator g, Eq. (34), interpolates the production over the sample of generating points $\{\hat{p}_{gj} | j \in \Lambda_g\}$ using convex coefficients $\beta_{gj,t}$ as computed in Eqs. (35) and (36). Consequently, in the objective, we seek the best interpolation for the cost functions. Though in general, we cannot foresee in which direction GLP errs, however, in the convex case, GLP overestimates the cost functions.

4.1 Illustration of GLP

Let us illustrate GLP on the following small quadratic problem:

$$f(x) = \min_{x_1, x_2} \left\{ \frac{1}{2} x_1^2 + x_2^2 - 2x_1 - 6x_2 \right\}$$
$$\text{S.t. } x_1 + x_2 \leq 1.75$$
$$-x_1 + 2x_2 \leq 2.2$$
$$2x_1 + x_2 \leq 4.7$$

Solving this problem with Cplex 12.5.1.0, the optimal solution is $x_1^* = 0.43333$, $x_2^* = 1.31667$, and the optimal value is $f^*(x) = -6.939167$. Table 1 shows the implementation of GLP for different sample points. GLP quickly converges to the true optimal solution as the grid points is slightly densified. Also observe that the function value is always overestimated.

Table 1: Example of GLP implementation

$\hat{x}_1$	$\hat{f}(x_1)$	$\hat{x}_2$	$\hat{f}(x_2)$
3	-1.5	7	7
31	418.5	21	315
x_1^*	x_2^*	$\hat{f}^*(x)$	True $f(x)$
–	–	–	–
$\hat{x}_1$	$\hat{f}(x_1)$	$\hat{x}_2$	$\hat{f}(x_2)$
3	-1.5	7	7
31	418.5	21	315
0	0.0	0	0.0
x_1^*	x_2^*	$\hat{f}^*(x)$	True $f(x)$
1.75	0.0	- 0.875	-1.96875
$\hat{x}_1$	$\hat{f}(x_1)$	$\hat{x}_2$	$\hat{f}(x_2)$
3	-1.5	7	7
31	418.5	21	315
0	0.0	0	0.0
9	22.5	5	-5
x_1^*	x_2^*	$\hat{f}^*(x)$	True $f(x)$
0.43333	1.31667	-1.5333	-6.93917
$\hat{x}_1$	$\hat{f}(x_1)$	$\hat{x}_2$	$\hat{f}(x_2)$
3	-1.5	7	7
31	418.5	21	315
0	0.0	0	0.0
9	22.5	5	-5
0.5	-0.875	1.5	-6.75
1.75	-1.96875	1.25	-5.9375
x_1^*	x_2^*	$\hat{f}^*(x)$	True $f(x)$
0.43333	1.31667	-6.5	-6.93917

5. Approximation of the cost-to-go function

Problem (32–36) cannot be solved for all state values (s_t, p_{t-1}, w_{t-1}). A common approach is to discretize s_t, p_{t-1}, and $\tilde{W}_{t-1}$ spaces into "partial grids", and to solve the problem over the Cartesian product of these grids. For instance, suppose we have a network with five wind turbines, five storage devices and five generators. In each period, if we discretize the wind turbine outputs, the production, and the storage into ten level each, without any consideration regarding the computational effort to optimize and to take the expectation,

we should solve the problem over a $10^5 \times 10^5 \times 10^5 = 10^{15}$-level grid, which is impossible in practice. Since its inception, dynamic programming (DP) has been limited to small instance of problems because of the "curse of dimensionality" as coined by its author, Bellman. The computation burden increases exponentially as the number of states increases.

These observations may have significant impact on the practicality of SDP as a solution method for real world problems. This is reinforced by the fact that in certain circumstances, the problem is to be solved periodically within constrained time frame. For instance, in the case of the unit commitment/economic dispatch problem, the operator of the system may need to adjust generating decisions based on signals from the market, or may have to resolve the problem periodically as a result of data changes or updated forecast as he/she gets to receive new observations for the random variables.

The practical limitation of DP paved the way for approximate dynamic programming (ADP) schemes. The purpose is usually to strike a balance between solution time and reasonable performance of the prescribed policy by replacing function $\tilde{F}_t(s_t, p_{t-1}, w_{t-1})$ with some approximation $\hat{F}_t(s_t, p_{t-1}, w_{t-1})$. Though he did not refer to as ADP, Bellman was the first to propose approximations of what he called the *functional equation* (e.g., here Eq. (17)). Using Lagrange relaxation technique, [7] (see also [9]) discusses successive approximations (SA) of the functional equation by partitioning the computation of the original sequence of functions into the computation of a sequence of functions of fewer state variables. [27] provide a detailed algorithm for the SA technique, and settle conditions for the convergence to true optimal solution. Examples of applications of such technique to power production are available in [64, 58, 61].

As indicated earlier, any DP implementation resorts to discretization of the state space. One straightforward approximation of the cost-to-go function is to select, in each period, a small sample of states instead of a dense grid, and solve problem (32–36) for each point of the sample, assuming that the minimization as well as the expectation computation can be carried out efficiently. Then the cost-to-go function may be approximated for any out-of-grid point, for instance by interpolation (e.g., linear, multilinear) of the neighboring grid states (see [56, 22, 24, 42]). Based on concavity assumptions on the cost-to-go function and assuming that the state space is a hyperrectangle, [62] propose a simplicial approximation scheme guided by local estimations of the approximation error. However, the complexity of hypercube decomposition limits the scope of the method.

Several approximation schemes fall within the broader class of *parametric approximation*, whereby the approximation may be denoted by $\hat{F}_t(s_t, p_{t-1}, w_{t-1}, \delta)$; δ is a set of parameters or weights, which usually are to be determined. For instance, in [8] such parameters are Legendre polynomial coefficients. Other polynomial types traditionally used to approximate the cost-to-go function include orthogonal, Chebyshev, spline, Hermite polynomials. For more account on polynomial approximations, see [22, 46, 20, 57].

[42] perform interpolations for out-of-grid states within hypercubes using weighted sums of the cost-to-go function evaluations as well as derivatives (first and second order) of the function at the vertices. The weights are defined as polynomials in the state variables. This multi-dimensional interpolation approach stems from previous works by [25, 26], and is an extension of the "gradient dynamic programming" scheme by [16].

Combining ideas from various fields such as *neural networks*, *artificial intelligence*, cognitive sciences, and so on, *reinforcement learning* (or *neuro-dynamic programming*) itera-

tively constructs approximations to the optimal cost-to-go function, or its expected value [12] through simulations. This methodology is applied by [23, 2, 59, 34], to power control decision problems.

Though there exist several reinforcement learning techniques, the most popular is the *Q-learning* algorithm, which in contrast with DP, computes the cost-to-go function for a set of randomly selected decisions considering only "visited states". For further details on contextual applications of this algorithm to power system problems, see [60, 36, 29, 55].

Reinforcement learning (neuro-dynamic programming) is built around *policy iteration*, and *value iteration*, two widespread DP algorithms [12]. The former algorithm iteratively alternates between *policy evaluation* and *policy improvement*, until no further improvement can be achieved (see [45, 44, 3]). The latter algorithm is a later name in the DP literature for successive approximation (see [53, 3]).

A complete review of ADP methods is beyond the scope of this paper.

6. Approximation through stochastic dual dynamic programming

Stochastic dual dynamic programming (SDDP) originated from the work by [41, 40, 39], and may be viewed as a variant of ADP techniques [50], where the cost-to-go function is approximated trough sampling and temporal decomposition, alternating between forward and backward steps.

Unlike other ADP methodologies, in SDDP, the state space is not discretized, but rather sampled. Then the cost-to-go function is approximated in the neighboring of the sampled states by supporting hyperplanes akin to *Benders' cuts* (see [10]). As a consequence, in principle, SDDP cannot be efficiently implemented for non-convex problems. Several extensions have been proposed to handle non-convexity. For instance, [18, 14] replace the non-convex productions functions with piecewise functions, and *McCormick envelopes*, respectively.

Let $V_{t+1}(s_{t+1}, p_t, \tilde{W}_t) = \mathbb{E}_{\tilde{W}_t|w_{t-1}}[\tilde{F}_{t+1}(s_{t+1}, p_t, \tilde{W}_t)]$. In each period, SDDP iteratively constructs an "outer" approximation to function $V_{t+1}(s_{t+1}, p_t, \tilde{W}_t)$,

$$\hat{V}_{t+1}(s_{t+1}, p_t, \tilde{W}_t) := \max_i \{H^i_{t+1}(s_{t+1}, p_t, \tilde{W}_t) | i \in I\},$$

where $H^i_{t+1}(s_{t+1}, p_t, \tilde{W}_t)$ is a hyperplane constructed using primal and dual information provided by the optimal solution. $\hat{V}_{t+1}(s_{t+1}, p_t, \tilde{W}_t)$ is therefore a lower bound on the true expected cost-to-go $V_t(s_t, p_{t-1}, \tilde{W}_{t-1})$. Figure 1 illustrates such an approximation for an hypothetical case.

Recall that for any function f of three variables x, y, and z, the first-order Taylor approximation, $\hat{f}$, in the neighboring of the point (x_0, y_0, z_0) is given by $\hat{f}(x, y, z) = f(x_0, y_0, z_0) + g'_x(x - x_0) + g'_y(y - y_0) + g'_z(z - z_0)$, which reduces to $\hat{f}(x, y, z) = c + g'_x x + g'_y y + g'_z z$, where $g_x \in \partial_x f(x_0, y_0, z_0) (g_y \in \partial_y f(x_0, y_0, z_0), \text{resp. } g_z \in \partial_z f(x_0, y_0, z_0))$ is a partial subgradient vector of f at $x (y, \text{resp. } z)$, and

$$c = f(x_0, y_0, z_0) - g'_x x_0 - g'_y y_0 - g'_z z_0. \tag{37}$$

For any vector a, g_a is a column vector, and g'_a is its transpose. We will henceforth drop the transposition operator for simplicity.

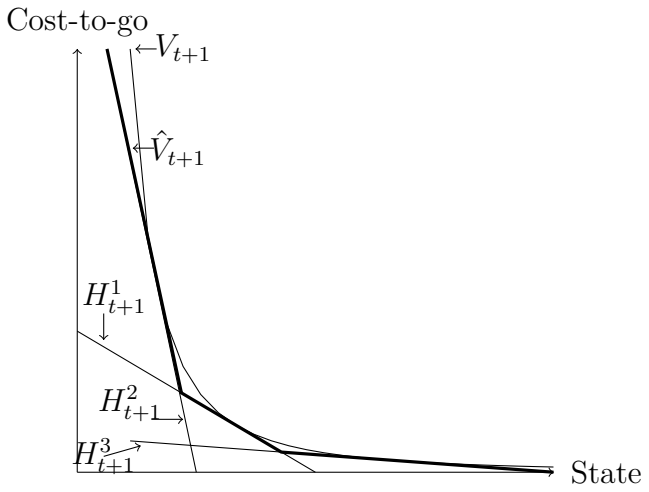

Figure 1: Illustration of the cost-to-go function approximation by SDDP

For a given set of "trial points" $\Theta_{t+1} = \left\{\left(s^i_{t+1}, p^i_t, w^i_t\right) | 1 \leq i \leq I\right\}$, we then have

$$H^i_{t+1}(s_{t+1}, p_t, w_t) = \tilde{c}^i_{t+1} + \tilde{g}_{s^i_{t+1}} s_{t+1} + \tilde{g}_{p^i_t} p_t + \tilde{g}_{w^i_t} w_t,\ 1 \leq i \leq I.$$

$\tilde{c}$ and $\tilde{g}$ are expected values with respect to the random variable $\tilde{W}_{t+1}$. It is readily verified that

$$\hat{V}_{t+1}(s_{t+1}, p_t, w_t) := \min\left\{\rho_{t+1} | \rho_{t+1} \geq \tilde{c}^i_{t+1} + \tilde{g}_{s^i_{t+1}} s_{t+1} + \tilde{g}_{p^i_t} p_t + \tilde{g}_{w^i_t} w_t,\ 1 \leq i \leq I\right\}.$$

As a result, an approximation to function $\tilde{F}_t(s_t, p_{t-1}, w_{t-1})$ reads:

$$\hat{F}_t(s_t, p_{t-1}, w_{t-1}) = \min_{P_t, \Delta_t, \kappa_t, s_{t+1}, \rho_{t+1}, \beta_t} \left\{ \sum_{g=1}^{|G|} h_t(\hat{p}, \Delta_t, \kappa_t) + \rho_{t+1} \right\} \tag{38}$$

$$\text{S.t.} (32) - (36) \tag{39}$$

$$\rho_{t+1} - \tilde{g}_{s^i_{t+1}} s_{t+1} - \tilde{g}_{p^i_t} p_t \geq \tilde{c}^i_{t+1} + \tilde{g}_{w^i_t} w_t, \qquad 1 \leq i \leq I, \tag{40}$$

where $h_t(\hat{p}, \Delta_t, \kappa_t) = \sum_{j \in \Lambda_g} \beta_{gj,t} CP_g(\hat{p}_g) + \sum_{j=1}^{|S|} CS_{jt}(\Delta_{jt}) + M \sum_{n=1}^{N} (\kappa^+_{nt} + \kappa^-_{nt})$.

6.1 Building the approximations

In period t, we solve problem (32–36) for a discrete set of state values (s_t, p_{t-1}, w_{t-1}) using an approximation, $\hat{V}_{t+1}(s_{t+1}, p_t, w_t)$, of the expected cost-to-go $V_{t+1}(s_{t+1}, p_t, w_t)$, defined as a set of I additional constraints to the problem (see problem (38–40)). These constraints are built using cut parameters ($\tilde{g}_{s_{t+1}}, \tilde{g}_{p_t}, \tilde{g}_{w_t}$, and $\tilde{c}_{t+1}$) computed in period $t+1$. To account for the uncertainty on the wind turbine outputs, suppose that in period $t+1$, we have considered a finite discrete set of realizations $\Omega_{t+1} = \left\{w^j_{t+1} | 1 \leq j \leq J\right\}$. Therefore, in period $t+1$, for each state value (s_{t+1}, p_t, w_t), the problem is solved for each $w_{t+1} \in \Omega_{t+1}$. For each such value of the random process, we have the following storage evolution equation:

$$s_{t+2} - \Delta^+_{t+1} + \Delta^-_{t+1} = \alpha s_{t+1}.$$

Let $\pi^j_{s,t+1}$ be the vector of dual multipliers associated with each such constraint. An expected partial subgradient $\tilde{g}_{s_{t+1}}$ may then be computed as

$$\tilde{g}_{s_{t+1}} = \alpha \sum_{j=1}^{J} \omega_j \pi^j_{s,t+1},$$

where ω_j is the probability of the observation w^j_{t+1}, and $\sum_{j=1}^{J} \omega_j = 1$.

In period $t+1$, for each generator $g \in G$, for each $w_{t+1} \in \Omega_{t+1}$, we also have

$$p_{g,t+1} \geq \max\{p_{gt} - \underline{\lambda}_g, \underline{p}_g\},$$

$$p_{g,t+1} \leq \min\{p_{gt} + \bar{\lambda}_g, \bar{p}_g\}.$$

Let π^j_{gl} and π^j_{gu} be the dual prices associated with these constraints, respectively. We know from duality theory (theorem of complementary slackness) that those multipliers are respectively non-null, if and only if the constraints are binding. Consequently, let

$$g^j_{gl} = \begin{cases} \pi^j_{gl} & \text{if } p^{j*}_{g,t+1} = p_{gt} - \underline{\lambda}_g, \\ 0 & \text{otherwise,} \end{cases}$$

where $p^{j*}_{g,t+1}$ is the optimal generation of the generator g. Similarly, let

$$g^j_{gu} = \begin{cases} \pi^j_{gu} & \text{if } p^{j*}_{g,t+1} = p_{gt} + \bar{\lambda}_g, \\ 0 & \text{otherwise.} \end{cases}$$

We then take

$$\tilde{g}_{p_t} = \sum_{j=1}^{J} \omega_j (g^j_l + g^j_u).$$

The generation of wind turbines located in the same region are correlated because of similar environmental conditions (wind speed). Consequently, suppose the wind generation is modeled as a lag-p multivariate autoregressive process:

$$w_{t+1} = \mu + \sum_{i=0}^{p-1} \Phi_j (w_{t-i} - \mu) + \tilde{\epsilon}_{t+1}, \tag{41}$$

where $\Phi_j, 0 \leq i \leq p-1$, is an $M \times M$ matrix of lag- $i+1$ autoregressive coefficients, μ is the mean vector of the process, $\tilde{\epsilon}_{t+1}$ is a vector of innovations. For a given state value (s_{t+1}, p_t, w_t), for each value w^j_{t+1} of the random process, and for each bus where a wind farm $m \in M$ is located, we have the power balance equation

$$\sum_{k=1}^{N_m} p_{k,t+1} + \sum_{l \in I_m} e_{l,t+1} - \gamma_{m,t+1} - \kappa^+_{m,t+1} + \kappa^-_{m,t+1} - \sum_{l \in O_m} e_{l,t+1} = \hat{D}_{m,t+1} - w^j_{m,t+1}. \tag{42}$$

Let $\pi^j_{wd,t+1}$ be the vector of dual multipliers associated with these constraints. In addition, suppose the expected cost-to-go function $\hat{V}_{t+2}(s_{t+2}, p_{t+1}, w_{t+1})$ is approximated through the following K cuts:

$$\rho_{t+2} - \tilde{g}_{s^k_{t+2}} s_{t+2} - \tilde{g}_{p^k_{t+1}} p_{t+1} \geq \tilde{c}^k_{t+2} + \tilde{g}_{w^k_{t+1}} w^j_{t+1}, 1 \leq k \leq K, \tag{43}$$

(see (38-40)). Let $\pi^{jk}_{wc,t+1}$ be the dual multiplier associated with inequality k. Using the dual prices associated with Eq. (42) and inequalities (43), respectively, and Eqs. (41), we compute a partial subgradient as

$$g_{w_t^j} = \Phi_1 \left(-\pi^j_{wd,t+1} + \sum_{k=1}^{K} \pi^{jk}_{wc,t+1} \tilde{g}_{w^k_{t+1}} \right).$$

Above equation follows from the chain rule $\frac{\partial \hat{F}_{t+1}}{\partial w_t} = \frac{\partial w_{t+1}}{\partial w_t} \frac{\partial \hat{F}_{t+1}}{\partial w_{t+1}}$. Thus, an expected partial subgradient $\tilde{g}_{w_t}$ is given by

$$\tilde{g}_{w_t} = \sum_{j=1}^{J} \omega_j g_{w_t^j}.$$

Lasly, from Eq. (37), we see that

$$\tilde{c}_{t+1} = \sum_{j=1}^{J} \omega_j \hat{F}_{t+1}(s_{t+1}, p_t, w_t) - \tilde{g}_{s_{t+2}} s_{t+2} - \tilde{g}_{p_{t+1}} p_{t+1} - \tilde{g}_{w_{t+1}} w_{t+1},$$

where $(s_{t+2}, p_{t+1}, w_{t+1})$ is the state observed in $t+2$ when (s_{t+1}, p_t, w_t) was observed in $t+1$.

Contrary to DP, where in each period we use a look-up table to keep track of the cost-to-go (we keep in memory the values of the discrete states as well as the corresponding cost-to-go), in SDDP, in each period, we only keep in memory the appropriate Lagrange multipliers and the expected optimal value $\hat{F}^*_t$, which are passed to the previous period.

The purpose of the backward recursion phase is to construct the approximations of the cost-to-go function, available in each period and in each iteration in the form of supporting hyperplanes to the true function. In each period (for each iteration), suppose we have a set of sampled states $\Theta_t = \left\{ \left(s^i_t, p^i_{t-1}, w^i_{t-1} \right) | 1 \leq i \leq I \right\}$. For each state vector, suppose we sample J vectors of wind generation $w^{ij}_t, 1 \leq j \leq J$, each with probability $\omega^j_i = 1/J$. Then for each wind output value, we solve the minimization problem (38-40). From the J minimization problems we retrieve the appropriate multipliers to compute the expected value of the parameters $\tilde{g}_{s^i_t}, \tilde{g}_{p^i_{t-1}}, \tilde{g}_{w^i_t}$, and $\tilde{c}^i_t$. These parameters are used to construct one supporting hyperplane $H^i_t(s_t, p_{t-1}, w_{t-1})$ for the minimization problem in period $t-1$. Thus, at the end of the recursion, we shall construct I supporting hyperplanes to the expectation of the cost-to-go function to pass to period $t-1$.

As initial conditions, we set $\Theta_{T+1} = \emptyset$, $\tilde{g}_{s_{T+1}} = 0, \tilde{g}_{p_T} = 0, \tilde{g}_{w_T} = 0$, and $\tilde{c}_{T+1} = 0$. A typical backward iteration is summarized in Table 2.

6.2 Assessing the quality of the approximations

In each iteration of the algorithm, we construct approximations to the true expected cost-to-go functions through backward recursions. As for any iterative scheme, after each iteration, we need to assess the quality of the approximations, and decide whether or not we can stop the algorithm. In SDDP, such an assessment step is carried out through forward simulations.

For a given initial state (s_1, p_0, w_0), suppose we simulate L series of wind generation vectors $w^l_t, 1 \leq t \leq T, 1 \leq l \leq L$, with probability $\omega_l = 1/L$ each. Solving problem (38-40) for the first period using the last approximation to the expected cost-to-go function constructed

Table 2: Backward induction procedure

For $t = T, T-1, \cdots, 1$
  Sample state variables $(s_t^i, p_{t-1}^i, w_{t-1}^i), 1 \leq i \leq I$.
  For $i = 1, \ldots, I$
    Sample J wind generation vectors w_t^{ij}, $1 \leq j \leq J$.
    For $j = 1, \ldots, J$
      Solve the minimization problem (38-40).
      Store the values of the appropriate multipliers and of $\hat{F}_t^*$.
    End
    Construct one supporting hyperplane $H_t^i(\cdot)$ for the problem in period $t-1$.
  End
End

in the backward step as well as the simulated wind generations, we obtain a lower bound on the true optimal cost (1–16). Indeed, as we indicated earlier, by construction, in the case of convexity, we have $\hat{F}_2(s_2, p_1, w_1) \leq \tilde{F}_2(s_2, p_1, w_1)$, which implies $\mathbb{E}_{\tilde{W}_2|w_1}\left[\hat{F}_2(s_2, p_1, w_1)\right] \leq \mathbb{E}_{\tilde{W}_2|w_1}\left[\tilde{F}_2(s_2, p_1, w_1)\right]$. It then follows that $\hat{F}_1(s_1, p_0, w_0) \leq \tilde{F}_1(s_1, p_0, w_0)$. Denote this lower bound by $\underline{F}$.

The decisions obtained from solving the problem in the first period are feasible, but not necessary optimal. By similarly computing the "immediate" decisions (production and storage variation) in each period $t, t = 2, \ldots, T$, using the cost-to-go approximations constructed in the backward recursion phase, we obtain a series of feasible, but not necessarily optimal management decisions. Thus, the sum of the immediate cost of the decisions so calculated may be used as an estimated upper bound on the true optimal total cost over the T planning span. Denoting this estimated upper bound by $\overline{F}$, we have

$$\overline{F} = \sum_{l=1}^{L} \sum_{t=1}^{T} h_t(p_t^l, \Delta_t^l, \kappa_t^l)\omega_l,$$

where $h_t(\cdot)$ is the immediate cost in period t.

Let $f_l = \sum_{t=1}^{T} h_t(p_t^l, \Delta_t^l, \kappa_t^l), 1 \leq l \leq L$. The uncertainty around the estimation of the upper bound may be estimated at

$$\sigma_{\bar{F}} = \sqrt{\frac{\sum_{l=1}^{l}(f_l - \overline{F})^2}{L-1}}.$$

This estimation may be used to construct a $(1-\alpha)\%$ confidence interval around the true value of the upper bound. Usually, if $\underline{F}$ is within the confidence interval, the algorithm is stopped assuming that the approximation is statistically accurate, otherwise, another backward approximation step is performed by adding new supporting hyperplanes (around the state values sampled during the simulation) to tighten the bounds. Drawbacks of such stopping criterion are discussed by [49]. [4, 43] instead use a fixed number of iterations as termination criterion. Also observe that in each iteration of the forward simulation, we

obtain l series of optimal trajectories for the storage levels as well as the generations. These may be used as sampled states for the next backward recursion step. Table 3 summarizes the forward procedure.

Table 3: Forward simulation procedure

Select an initial state (s_1, p_0, w_0).
Simulate L series of wind generations each of length T.
For $l = 1, \ldots, L$
For $t = 1, \ldots, T$
Solve the minimization problem (38-40).
Store the optimal decisions $s_{t+1}^{l^*}$ and $p_t^{l^*}$ and the wind value w_t^l.
End
End

7. Illustration and comparison with classical DP

We used the IEEE 9-bus configuration (Figure 2) to illustrate and compare SDDP algorithm with classical DP. This network comprises three conventional generators for a total capacity of 820 MW, and nine transmission lines. We used NYISO scaled average hourly load data for January 2016 [37]. We located one wind farm and one storage facility at bus 5. Wind data were obtained from the website of the National Renewable Energy Laboratory [28]. We utilized a 15% wind integration scenario and assumed 30% wind capacity factor.

Both SDDP and DP algorithms were implemented in Python 2.7.10, and the GLP problems were solved with Cplex 12.5.1.0. At each iteration of the SDDP backward pass, we simulated 25 vectors of wind outputs to compute the expected values of the cost-to-go functions as well as the expected values of the partial subgradients $(\tilde{g}_{s_t}, \tilde{g}_{p_{t-1}}, \tilde{g}_{w_{t-1}})$, and the expected values of the intercepts $(\tilde{c}_t^i)$. We also used 25 series of state values in each iteration of the forward pass. As a result, in each iteration of the backward pass, 25 news cuts were added to the GLP problems.

We considered a hypothetical battery with maximum capacity of 250 MWh. Maximum hourly charging or discharging was assumed to be 30% of the battery capacity, and the minimum level of the stored energy to be 20% of the battery capacity. For the sake of illustration, we supposed that the battery had 100% storage efficiency.

Figure 3 depicts the trajectories of the storage facility over a 24 hour span for two cases. In case (a) the stored energy is at the minimum allowable level in the beginning of the first hour; in case (b) the level of charge is at 60% in the beginning of the first hour. We also assumed zero-cost for varying the level of energy in this example. The net load (demand - wind output) is also depicted in each case.

In the first case, the battery is charged progressively in the beginning as the net load is low, then is steady as of the first peak hour, and up to the second peak hour. Then the battery is discharged progressively to its minimum level, since the model considered no

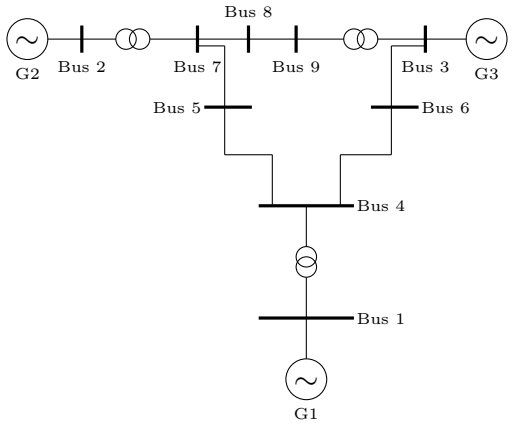

Figure 2: 9-bus system configuration

"terminal value" for the storage. In the second situation, the level of charge is steady until the first peak hour, then the battery is depleted until the end of the horizon.

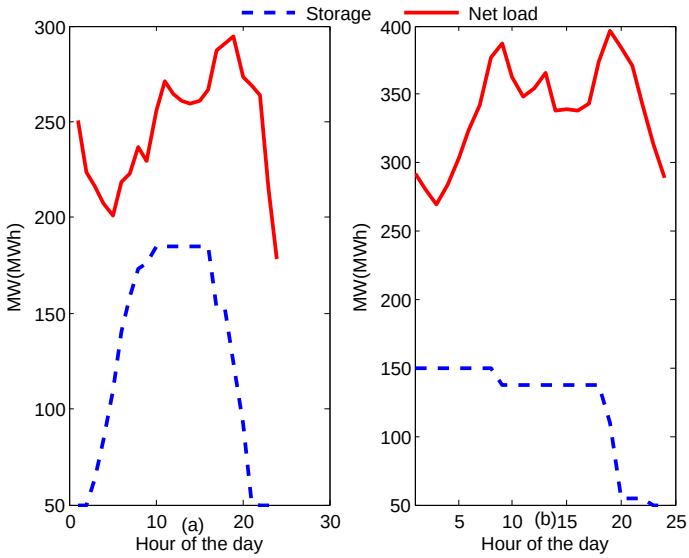

Figure 3: Example of storage trajectory when charging (discharging) cost is low

Figure 4 illustrates a situation where varying the level of charge of the battery is very expensive as compared to the cost of operating the conventional units. Since utilizing the battery is very costly, the stored energy is used only in hours where the conventional generators cannot meet the demand.

How would SDDP compare with classical DP? In general, this would difficult to answer due to the prohibitiveness of the computational burden of DP. We compared the performance of both algorithms on the 9-bus network example based on two criteria, namely (i) the CPU time, and (ii) the solution performance (total cost). We first approximated the value functions with both SDDP and DP. For DP, in each hour, the storage, and the generators outputs were discretized each into six levels, and the wind farm output into seven levels.

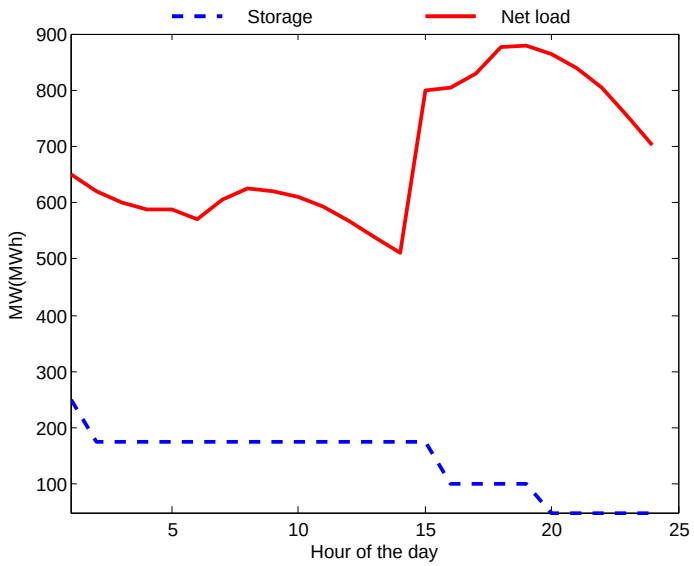

Figure 4: Example of storage trajectory when charging (discharging) cost is high

This was found to be the finest grids to solve the problems in reasonable time. Since we assumed the wind process to be Markovian, in each hour, we needed to consider all the transitions from the previous wind values (e.g. 7×7 transitions). As a result, in each hour, the DP problem was solved over a grid of $6 \times 6^3 \times 7^2$ levels (the network comprising three conventional generators), which resulted in 63 504 evaluations of the cost-to-go function in each hour. Therefore, 1 524 096 evaluations were performed over the 24-hour horizon.

Following the construction of the cost-to-go functions with both methods, we simulated the operation of the network for two simulation runs, each comprising one hundred 24-hour wind scenarios. With DP, in both run, the cost-to-go functions were constructed over the same state space. With SDDP, in the first run, we performed only four iterations to approximate the value functions. Since in each hour we approximated the expected cost-to-go using 25 supporting hyperplanes (state values) and used 20 wind values to compute the expectations, with SDDP, the total number of function evaluations then were 500, for a total of 12 000 function evaluations over the 24 hour horizon. Therefore, overall 48 0000 function evaluations were performed in the first simulation. Those numbers do not include the SDDP forward step - to assess the quality of the approximation-, since this is very fast. In the second simulation, we tried to improve the quality of the approximation by carrying out ten iterations of SDDP. Consequently, 120 000 function evaluations were performed in that run.

Table 4 presents the CPU time for each method and each run. In the first case, the computation time for SDDP was about 4% of that of DP. In the second case, that proportion was about 15% as better approximations were sought.

Table 5 reports descriptive statistics on the performance (total cost) of each method for each run. In the first case, on average (over the 100 scenarios), using DP, the total cost was only improved by 0.71% as compared to SDDP. In the second case, on average,the total cost difference decreased to 0.21%. This suggests that, overall, SDDP allowed for a fair trade-off

Table 4: Comparison of the CPU time in seconds: SDDP and DP

Method	Run 1	Run 2
SDDP	499.35	1 876.22
DP	13 038.16	12 726.88

between solution time and accuracy.

Table 5: Comparison of solution cost: SDDP and DP

Method	Run 1			
	Min	Max	Mean	Standard deviation
SDDP	78 622.29	162 215.38	126 872.83	21 812.53
DP	75 392.76	161 872.07	125 968.26	22 922.10
Method	Run 2			
	Min	Max	Mean	Standard deviation
SDDP	88 691.21	164 333.09	134 267.42	19 210.77
DP	85 882.05	164 333.09	133 981.13	19 657.52

Figure 5 depicts the average storage trajectory for both method and both simulation runs. In general, with DP, the level of charge is higher, but the overall pattern (time of charging and discharging) is similar.

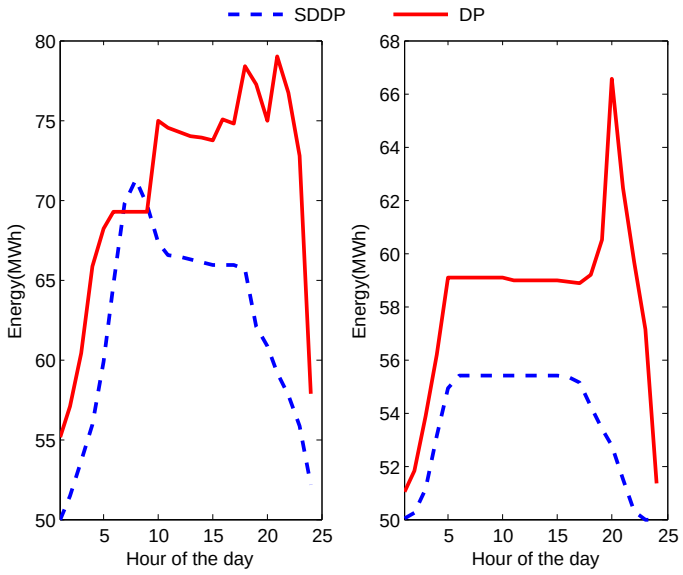

Figure 5: Mean storage trajectory obtained from SDDP and DP

8. Additional numerical tests

The ultimate goal of this research was to analyze the scalability of SDDP to larger networks. We tested the algorithm on different IEEE networks, which characteristics are presented in

Table 6. Each network was tested with different number of storage facilities and wind farms. Wind and load data were obtained from the sources as in Section 7.

Table 6: Test networks' characteristics

Case	# of buses	# of generators	# of trans. lines
1	30	6	41
2	57	7	80
3	89	12	210
4	118	54	186
5	300	69	411

Figure 6 shows an example of the mean trajectory over 100 simulations for a network composed of 118 buses, 5 storage facilities and one aggregated wind farm, as well as the net load. Each battery has specific storage, charging, and discharging efficiency, respectively. The three batteries that either have the highest storage or charging/discharging efficiency are used to contribute to meet the load, whereas the energy level of the other two batteries (with either the lowest storage or charging/discharging efficiency) is steady over the 24-hour span.

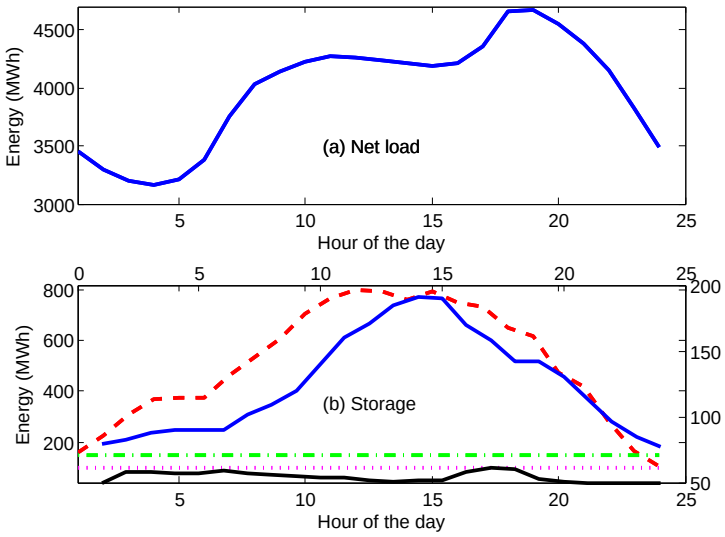

Figure 6: Mean storage trajectory for the 118-bus network with five storage units and one wind farm

Table 7 reports the computation time for each network, and equipped with different number of storage units (|S|) and wind farms (|M|). We varied the number of storage facilities from one to twenty, and the number of wind farms from one to ten. The computation time varied between: (i) 1 230 seconds and 1 583 seconds in the case of the 30- bus network, (ii) 1 388 seconds and 1 598 seconds for the 57-bus network, (iii) 1 570 and 1 737 seconds for the 57-bus network, (iv) 2 180 seconds and 2 454 seconds in the case of the 118-bus network, and, (v) 4 159.16 seconds and 5 036.37 seconds for the largest network (300 buses).

Table 7: Computation time in seconds

# buses	|S|	|M|	CPU time	# buses	|S|	|M|	CPU time
30	1	1	1 229.80	118	1	1	2 399.35
30	5	1	1 582.67	118	5	1	2 444.99
30	5	5	1 323.88	118	5	5	2 453.89
57	1	1	1 388.09	118	10	5	2 179.62
57	5	1	1 454.47	118	20	10	2 248.39
57	5	5	1 396.26	300	1	1	4 159.16
57	10	5	1 597.71	300	5	1	4 234.72
89	1	1	1 570.67	300	5	5	4 570.01
89	5	1	1 709.68	300	10	5	4 617.65
89	5	5	1 575.09	300	20	10	5 036.37
89	10	5	1 737.32				

9. Conclusions

This paper analyzed the operation of power networks comprising both conventional and wind generators in the presence of storage. We formulated the problem under the framework of dynamic programming, and used stochastic dual dynamic programming and generalized linear programming to approximate the problem. Such approximation schemes allowed to handle large dimension state space and to solve the problem in reasonable time as compared to classical dynamic programming.

References

[1] Chad Abbey and Géza Joós. A stochastic optimization approach to rating of energy storage systems in wind-diesel isolated grids. *IEEE Transactions on Power Systems*, 24(1):418–426, 2009.

[2] TP Imthias Ahamed, PS Nagendra Rao, and PS Sastry. A reinforcement learning approach to automatic generation control. *Electric power systems research*, 63(1):9–26, 2002.

[3] Roger N Anderson, Albert Boulanger, Warren B Powell, and Warren Scott. Adaptive stochastic control for the smart grid. *Proceedings of the IEEE*, 99(6):1098–1115, 2011.

[4] Tsvetan Asamov and Warren B Powell. Regularized decomposition of high-dimensional multistage stochastic programs with markov uncertainty. *arXiv preprint arXiv:1505.02227*, 2015.

[5] John P Barton and David G Infield. Energy storage and its use with intermittent renewable energy. *IEEE Transactions on Energy Conversion*, 19(2):441–448, 2004.

[6] John P Barton and David G Infield. A probabilistic method for calculating the usefulness of a store with finite energy capacity for smoothing electricity generation from wind and solar power. *Journal of Power Sources*, 162(2):943–948, 2006.

[7] Richard Bellman. Dynamic programming and lagrange multipliers. *Proceedings of the National Academy of Sciences of the United States of America*, 42(10):767, 1956.

[8] Richard Ernest Bellman. Some new techniques in the dynamic-programming solution of variational problems. 1957.

[9] Richard Ernest Bellman and Stuart E Dreyfus. Applied dynamic programming. 1962.

[10] Jacques F Benders. Partitioning procedures for solving mixed-variables programming problems. *Numerische mathematik*, 4(1):238–252, 1962.

[11] Dimitri P Bertsekas. *Dynamic programming and optimal control 3rd edition, volume I*. Belmont, MA: Athena Scientific, 2005.

[12] Dimitri P Bertsekas and John N Tsitsiklis. Neuro-dynamic programming: an overview. In *Decision and Control, 1995., Proceedings of the 34th IEEE Conference on*, volume 1, pages 560–564. IEEE, 1995.

[13] Mary Black and Goran Strbac. Value of storage in providing balancing services for electricity generation systems with high wind penetration. *Journal of Power Sources*, 162(2):949–953, 2005.

[14] Santiago Cerisola, Jesus M Latorre, and Andres Ramos. Stochastic dual dynamic programming applied to nonconvex hydrothermal models. *European Journal of Operational Research*, 218(3):687–697, 2012.

[15] James M Eyer, Joseph J Iannucci, and Garth P Corey. Energy storage benefits and market analysis handbook, a study for the doe energy storage systems program. *Sandia National Labratories*, 2004.

[16] Efi Foufoula-Georgiou and Peter K Kitanidis. Gradient dynamic programming for stochastic optimal control of multidimensional water resources systems. *Water Resources Research*, 24(8):1345–1359, 1988.

[17] Javier Garcia-Gonzalez, De la Muela, Rocío Moraga Ruiz, Luz Matres Santos, and Alicia Mateo González. Stochastic joint optimization of wind generation and pumped-storage units in an electricity market. *IEEE Transactions on Power Systems*, 23(2):460–468, 2008.

[18] Quentin Goor, R Kelman, and Amaury Tilmant. Optimal multipurpose-multireservoir operation model with variable productivity of hydropower plants. *Journal of Water Resources Planning and Management*, 137(3):258–267, 2010.

[19] Ioannis Hadjipaschalis, Andreas Poullikkas, and Venizelos Efthimiou. Overview of current and future energy storage technologies for electric power applications. *Renewable and Sustainable Energy Reviews*, 13(6):1513–1522, 2009.

[20] Richard Howitt, Siwa Msangi, Arnaud Reynaud, and Keith Knapp. Using polynomial approximations to solve stochastic dynamic programming problems: Or a 'betty crocker'approach to sdp. *University of California, Davis*, 2002.

[21] Hussein Ibrahim, Adrian Ilinca, and Jean Perron. Energy storage systems—characteristics and comparisons. *Renewable and Sustainable Energy Reviews*, 12(5):1221–1250, 2008.

[22] Sharon A Johnson, Jery R Stedinger, Christine A Shoemaker, Ying Li, and Jose Alberto Tejada-Guibert. Numerical solution of continuous-state dynamic programs using linear and spline interpolation. *Operations Research*, 41(3):484–500, 1993.

[23] Rajit Johri and Zoran Filipi. Self-learning neural controller for hybrid power management using neuro-dynamic programming. 2011.

[24] Michael P Keane and Kenneth I Wolpin. The solution and estimation of discrete choice dynamic programming models by simulation and interpolation: Monte carlo evidence. *The Review of Economics and Statistics*, pages 648–672, 1994.

[25] PK Kitanidis. Hermite interpolation on an n-dimensional rectangular grid, 1986.

[26] PK Kitanidis. A first-order approximation to stochastic optimal control of reservoirs. *Stochastic Hydrology and Hydraulics*, 1(3):169–184, 1987.

[27] AJ Korsak and RE Larson. A dynamic programming successive approximations technique with convergence proofs. *Automatica*, 6(2):253–260, 1970.

[28] National Renewable Energy Laboratory. `Available:http://www.nrel.gov/electricity/transmission/eastern_wind_dataset.html`.

[29] Jin-Hee Lee and John W Labadie. Stochastic optimization of multireservoir systems via reinforcement learning. *Water Resources Research*, 43(11), 2007.

[30] Steve Lindenberg. *20% Wind Energy By 2030: Increasing Wind Energy's Contribution to US Electricity Supply*. Diane Publishing, 2009.

[31] Jim McDowall. Integrating energy storage with wind power in weak electricity grids. *Journal of Power Sources*, 162(2):959–964, 2005.

[32] Peter Meibom, Rüdiger Barth, Bernhard Hasche, Heike Brand, Christoph Weber, and Mark O'Malley. Stochastic optimization model to study the operational impacts of high wind penetrations in ireland. *IEEE Transactions on Power Systems*, 26(3):1367–1379, 2011.

[33] Pedram Mokrian and Moff Stephen. A stochastic programming framework for the valuation of electricity storage. In *26th USAEE/IAEE North American Conference*, pages 24–27, 2006.

[34] James Momoh, Yi Zhang, et al. Unit commitment using adaptive dynamic programming. In *Intelligent Systems Application to Power Systems, 2005. Proceedings of the 13th International Conference on*, pages 4–pp. IEEE, 2005.

[35] Pedro S Moura and Anibal T De Almeida. The role of demand-side management in the grid integration of wind power. *Applied Energy*, 87(8):2581–2588, 2010.

[36] MB Naghibi-Sistani, MR Akbarzadeh-Tootoonchi, MH Javidi-Dashte Bayaz, and H Rajabi-Mashhadi. Application of q-learning with temperature variation for bidding strategies in market based power systems. *Energy Conversion and Management*, 47(11):1529–1538, 2006.

[37] New York Independent System Operator. Market and operations. `Available:http://www.nyiso.com/public/markets_operations/market_data/load_data/index.jsp`.

[38] Anthony Papavasiliou and Shmuel S Oren. Multiarea stochastic unit commitment for high wind penetration in a transmission constrained network. *Operations Research*, 61(3):578–592, 2013.

[39] Mario VF Pereira and Leontina MVG Pinto. Multi-stage stochastic optimization applied to energy planning. *Mathematical programming*, 52(1-3):359–375, 1991.

[40] MVF Pereira. Optimal stochastic operations scheduling of large hydroelectric systems. *International Journal of Electrical Power & Energy Systems*, 11(3):161–169, 1989.

[41] MVF Pereira and LMVG Pinto. Stochastic optimization of a multireservoir hydroelectric system: a decomposition approach. *Water resources research*, 21(6):779–792, 1985.

[42] C Russell Jr Philbrick and Peter K Kitanidis. Improved dynamic programming methods for optimal control of lumped-parameter stochastic systems. *Operations Research*, 49(3):398–412, 2001.

[43] Andrew B Philpott and Vitor L de Matos. Dynamic sampling algorithms for multi-stage stochastic programs with risk aversion. *European Journal of Operational Research*, 218(2):470–483, 2012.

[44] Qinru Qiu and Massoud Pedram. Dynamic power management based on continuous-time markov decision processes. In *Proceedings of the 36th annual ACM/IEEE Design Automation Conference*, pages 555–561. ACM, 1999.

[45] JP Quadrat, M Akian, and JP Chancelier. Dynamic programming complexity and application. In *Proceedings of the 27th Conference on Decision and Control*, pages 1551–1558, 1988.

[46] John Rust. Numerical dynamic programming in economics. *Handbook of computational economics*, 1:619–729, 1996.

[47] Susan Schoenung. Energy storage systems cost update, 2011.

[48] W Scott and Warren B Powell. Approximate dynamic programming for energy storage with new results on instrumental variables and projected bellman errors. *Submitted to Operations Research (Under Review)*, 2012.

[49] Alexander Shapiro. Analysis of stochastic dual dynamic programming method. *European Journal of Operational Research*, 209(1):63–72, 2011.

[50] Alexander Shapiro, Wajdi Tekaya, Joari Paulo da Costa, and Murilo Pereira Soares. Risk neutral and risk averse stochastic dual dynamic programming method. *European journal of operational research*, 224(2):375–391, 2013.

[51] Jeremy F Shapiro. *Mathematical Programming: Structures and Algorithms*. Wiley New York, 1979.

[52] J Charles Smith, Michael R Milligan, Edgar DeMeo, Brian Parsons, et al. Utility wind integration and operating impact state of the art. *IEEE Transactions on Power Systems*, 22(3):900–908, 2007.

[53] Haili Song, Chen-Ching Liu, Jacques Lawarrée, and Robert W Dahlgren. Optimal electricity supply bidding by markov decision process. *IEEE Transactions on Power Systems*, 15(2):618–624, 2000.

[54] Samir Succar, David C Denkenberger, and Robert H Williams. Optimization of specific rating for wind turbine arrays coupled to compressed air energy storage. *Applied Energy*, 96:222–234, 2012.

[55] Ying Tan, Wei Liu, and Qinru Qiu. Adaptive power management using reinforcement learning. In *Proceedings of the 2009 International Conference on Computer-Aided Design*, pages 461–467. ACM, 2009.

[56] J Alberto Tejada-Guibert, Sharon A Johnson, and Jery R Stedinger. Comparison of two approaches for implementing multireservoir operating policies derived using stochastic dynamic programming. *Water resources research*, 29(12):3969–3980, 1993.

[57] Huseyin Topaloglu and Warren B Powell. Dynamic-programming approximations for stochastic time-staged integer multicommodity-flow problems. *INFORMS Journal on Computing*, 18(1):31–42, 2006.

[58] Andre Turgeon. Optimal operation of multireservoir power systems with stochastic inflows. *Water resources research*, 16(2):275–283, 1980.

[59] John G Vlachogiannis and Nikos D Hatziargyriou. Reinforcement learning for reactive power control. *Power Systems, IEEE Transactions on*, 19(3):1317–1325, 2004.

[60] Gaofeng Xiong, Tomonori Hashiyama, and Shigeru Okuma. An electricity supplier bidding strategy through q-learning. In *Power Engineering Society Summer Meeting, 2002 IEEE*, volume 3, pages 1516–1521. IEEE, 2002.

[61] Jin-Shyr Yang and Nanming Chen. Short term hydrothermal coordination using multi-pass dynamic programming. *Power Systems, IEEE Transactions on*, 4(3):1050–1056, 1989.

[62] Luckny Zéphyr, Pascal Lang, and Bernard F Lamond. Controlled approximation of the value function in stochastic dynamic programming for multi-reservoir systems. *Computational Management Science*, 12(4):539–557, 2015.

[63] Yangfang Zhou, Alan Andrew Scheller-Wolf, Nicola Secomandi, and Stephen Smith. Managing wind-based electricity generation in the presence of storage and transmission capacity. *Available at SSRN 1962414*, 2014.

[64] HH Zurn and VH Quintana. Generator maintenance scheduling via successive approximations dynamic programming. *Power Apparatus and Systems, IEEE Transactions on*, 94(2):665–671, 1975.